FIFTH EDITION

PERSPECTIVES ON LEARNING

D. C. PHILLIPS
JONAS F. SOLTIS

Published by Teachers College Press, 1234 Amsterdam Avenue, New York, NY 10027

Library of Congress Cataloging-in-Publication Data

Phillips, D. C. (Denis Charles), 1938-
Perspectives on learning / D. C. Phillips, Jonas F. Soltis. — 5th ed.
p. cm. — (Thinking about education series)
Includes bibliographical references and index.
ISBN 978-0-8077-4983-8 (pbk. : alk. paper) 1. Learning—Philosophy. I. Soltis, Jonas F. II. Title.
LB1060.P48 2009
370.15'23--dc22

2009006147

ISBN 978-0-8077-4983-8 (paper)

Printed on acid-free paper

Manufactured in the United States of America

16 15 14 13 12 11 10 09 8 7 6 5 4 3 2 1

Contents

Acknowledgments ix

A Note to the Instructor xi

Chapter 1
INTRODUCTION 1
- Learning 1
- The Teacher's Responsibility 2
- The Variety of Theories 3
- An Objection: We Don't Need Theories, Just Common Sense 5
- The Plan of the Book 6

Chapter 2
CLASSICAL THEORIES 9
- Plato's Theory of Learning 10
- Case One 12
- The Lockean Atomistic Model 13
- Case Two 15
- A Critique 16
- Case Three 18

Chapter 3
BEHAVIORISM 21
- Classical Conditioning 23
- Operant Conditioning 25
- Case One 26
- B. F. Skinner 27
- Strengths and Weaknesses 28
- Case Two 30

Chapter 3

Behaviorism

The classical theories of learning discussed in the previous chapter were put forward before the field of psychology had separated off from philosophy as a discipline in its own right around the end of the nineteenth century. Plato and Locke were philosophers who had an interest in what we would now regard as psychological topics, but their method of investigating these topics was philosophical, not scientific. In fact, psychology only became a science—with the paraphernalia of experimentation, careful and systematic observation, measurement, and calculation—long after Locke's time. (The years around 1870 are often selected as marking the birth of modern experimental psychology.) The behaviorists worked after this watershed period; they were completely saturated with the scientific spirit—so much so that they may eventually have been misled, for arguably they held a very narrow view of the nature of science that did not, for example, accommodate unobservable entities or events (such as mental phenomena). But their work has thrown a great deal of light on certain types of learning, and some of the principles they formulated are still useful ingredients in the armament of the teacher.

After publication of Darwin's theory of evolution in 1859, humans were seen as being "biologically continuous" with the animal kingdom. Humans may be different in many ways from the other species of animals, but at least so far as biological equipment is concerned (physiology, biochemistry, anatomy, and so forth) it became recognized that there were great similarities. In light of this, it seemed reasonable to expect that by studying biological processes in animals some insight would be gained into similar processes in humans. The great advances that have since taken place in the medical sciences have justified this approach. It was not long before researchers interested in psychological phenomena started working along the same lines. There was a spate of activity in the closing years of the nineteenth century onwards directed at how animals learn, the nature of drives and instincts, problem solving, and the like. Behaviorism developed out of this background.

Researchers working with animals seem to have a serious disadvantage that is not present for those psychologists who are studying humans. You can ask people to introspect and report what is happening in their

minds—what they are thinking, what they are experiencing, what they remember, what they can see, how they feel—but there is not much point in asking a dog or cat or pigeon! Animal researchers, then, seem to lack this vital introspective source of data that is available to those who work on their fellow humans. It occurred to John B. Watson (1878–1958) that this supposed disadvantage was, in fact, very much an advantage.

The situation as Watson saw it in 1913 was that researchers working with humans had no reliable way of validating the introspective reports that people gave about what was happening in their minds. What is happening in your mind is accessible to one person only—you. But the essence of science is objectivity and replication; data have to be accessible to many researchers so that findings can be checked or reproduced by others. So Watson recommended that psychologists working with humans should use the same methods that animal researchers were using. The opening paragraph of his famous article in the *Psychological Review* in 1913, "Psychology as the Behaviorist Views It," was a vigorous call for a revolution, and it gave birth to behaviorism:

> Psychology as the behaviorist views it is a purely objective experimental branch of natural science. Its theoretical goal is the prediction and control of behavior. Introspection forms no essential part of its methods, nor is the scientific value of its data dependent upon the readiness with which they lend themselves to interpretation in terms of consciousness. The behaviorist, in his efforts to get a unitary scheme of animal response, recognizes no dividing line between man and brute.[1]

Every scientist must make assumptions in order to get his or her work started. One has to start somewhere. There has to be some basis from which the work will take off. Watson's paper is remarkable for the clarity and honesty with which it makes these assumptions or foundations explicit. But Watson's assumptions are tough ones. At first sight it looks extremely difficult to throw light on human learning, not by asking people about what they have learned, but by concentrating instead on their observable behavior. And it does not look too hopeful to try to understand how a person can learn something as complex as Einstein's theory, or debating techniques, by using principles and methods that were developed to study learning in pigeons and laboratory rats.

However, like Locke, the behaviorists assumed that humans were biologically "wired" or equipped so that they could interact with the environment, and profit from this interaction. But, unlike Locke and Plato, the behaviorists were not concerned about the origin of ideas or concepts. Plato's problem of how the acquisition of knowledge is possible unless something prior is known was not a problem for them—the behaviorists conceptualized matters differently. For the behaviorists, the issue was

not how new *knowledge* is acquired; instead it was: How is new *behavior* acquired? In other words, to the behaviorists learning was a process of expanding the behavioral repertoire, not a matter of expanding the ideas in the learner's mind. (Mind, after all, was a subjective and nonpublicly observable entity, and thus had to be avoided by science.) In terms of a classroom example, the behaviorist is interested not in how a pupil understands and learns Einstein's theory, but in how a pupil can be led to behave in such a way that he or she can do certain things (such as get the correct answer to problems, perform experiments, write down certain equations when asked by a teacher, and so on).

How, then, do people learn, in the behaviorist's sense of acquiring new behaviors? What scientifically observable factors are at work? The answers to these questions come in two parts, for there are two traditions within behaviorism: on one hand there is classical conditioning or stimulus substitution behaviorism, and on the other, operant conditioning or response reinforcement behaviorism.

Classical Conditioning

Around the beginning of the twentieth century the Russian physiologist Ivan Pavlov was studying the process of digestion in dogs, and by chance he discovered he was producing an interesting change in their behavior. (These days scientists have a technical name to make chance discoveries sound more impressive—they call it serendipity.) Pavlov noticed that whenever he fed the dogs, they started to produce saliva. It was produced even at the sight of food, just as you might "water at the mouth" when you see an ad for pizza or hamburgers on TV. Investigating further, he found that if he rang a bell at the same time that he fed the animals, after many repetitions just the sound (without the sight of the meal) led the dogs to salivate. The dogs had been conditioned.

It turns out that dogs, and indeed humans, have an inborn mechanism that stimulates the production of saliva when food is present (in popular language such a mechanism is often called a reflex action). The animal is "biologically wired" so that a certain stimulus (in this case, food) produces a specific response (salivation). Diagrammatically:

stimulus (food) ——▶ response (salivation)

or

S ——▶ R

What had happened in Pavlov's laboratory was that, in this natural reflex, a new stimulus (the bell) had become coupled with the natural stimulus

(the food), and eventually it was able by itself to produce the original or natural response (the salivation). In other words, the new or conditioned stimulus replaced the natural stimulus. Diagrammatically:

1. natural stimulus (food)——————▶ response (salivation)

2. natural stimulus (food)
plus——————————▶ response (salivation)
conditioned stimulus (bell)

and eventually,

3. conditioned stimulus (bell)————▶ response (salivation)

Watson was familiar with Pavlov's work, and he seized upon classical conditioning as the key mechanism underlying all human learning. Consider a homely example: How does a young child acquire the behavior of being difficult to put to bed in the evening? This is something the child has learned—but how? Who taught the child this undesirable way of behaving? The key is to find a natural reflex that produces the response we are considering (in this example, the difficult behavior at bedtime). Somehow the natural stimulus that produces this behavior has been replaced by a conditioned stimulus. Thus, a loud noise can produce fear and unsettled behavior—this is a natural or "wired in" reflex. Perhaps one night when the child was being put to bed, the parent slammed the door at the same time as putting out the light, and if the fright was severe enough the child might have become conditioned by this one event. In other words, the natural stimulus (the frightening noise) was coupled with the conditioned stimulus (the light being extinguished). The net result: whenever the light is turned off, the child becomes scared and unruly.

Watson saw these "built-in" behaviors everywhere (he sometimes called them built-in, rather than conditioned, behaviors because he wanted to stress for parents and teachers that they bore a great deal of responsibility for the things the child learned). He wrote:

> The main point to emphasize is that practically every responding organ of the body can be conditioned; and that this conditioning takes place not only throughout adult life but can and does take place daily from the moment of birth (in all probability before birth).... All of us are shot through with stimulus substitutions of one kind or another which we know nothing about until the behaviorist tries us out and tells us about them. [2]

Undoubtedly Watson was right. Every time a teacher shouts angrily at a class of pupils reading Shakespeare, a child might be adversely conditioned; the natural response to shouting and anger is to avoid whatever it was that the shouting was about. If the shouting or anger accompanied the reading of Shakespeare, then the Bard might in future be avoided like the plague. And if working under relaxed and quiet conditions naturally produces feelings of pleasure, and if mathematics is always done when these conditions prevail, then mathematics might become conditioned and produce feelings of pleasure. But to say that we are shot through with conditioned reflexes is not to say that our behavior consists *only* of conditioned reflexes. There are many cases of learning where accounts in terms of classical conditioning seem to be completely far-fetched. How can conditioned reflexes account, for example, for the learning of something as abstract as Einstein's theory, or even the learning of language?

All is not lost, for there is another mechanism to which the behaviorist can appeal.

Operant Conditioning

E. L. Thorndike (1874–1949) received his doctorate in 1898 for his study of the learning behavior of cats. Typically he would imprison a cat inside a box, which would have some simple release mechanism such as a chain or a lever that the cat could activate by touching. Upon release, the cat would be rewarded in some fashion—it would be given access to a piece of fish, for example. Thorndike recorded the time the cat took to escape over successive trials, and he graphed these to produce what is now known as a learning curve. It might be expected that on the first few trials the cat would be slow to escape, but soon it would "get the idea" and would learn to escape almost instantaneously. Instead, Thorndike found that there was a constant and gradual improvement; on each successive trial the cat would escape a little more rapidly, so that after many attempts (say, about a dozen) the cat would be getting out as soon as it was placed in the box.

As a result of this kind of work, Thorndike formulated several "laws of learning." His conception was that some nerve pathway had been established in the brain of the learning animal, so that when a particular stimulus was registered by the sense organs it became connected by this pathway to the organs that produced the response that had proven to be effective. His "law of exercise" was that the link between a given stimulus and response becomes stronger the more the pathway is activated; in other words, the more the behavior is practiced or exercised, the more strongly it will be established or "learned." His "law of effect" stated that if the response to a stimulus has a pleasing effect, then the probability would increase of the learner repeating that response when confronted with the same stimulus. (In terms of the cat, the first law means that the more times

the cat practices escaping from the box, the better it will be able to do it; the second law means that if the response—the escaping—has pleasant consequences, such as access to fish, then if the cat is placed in the same stimulus situation again, it is more likely to repeat the rewarded response.) Diagrammatically, the "law of effect" can be depicted:

(imprisonment in box)	(push lever)	(freedom or fish)
stimulus ———▶	response ———▶	reward or reinforcement, which makes the response more probable

Thorndike was quite sure that these laws applied to humans as much as to simpler mammals. He wrote in 1913:

> These simple, semi-mechanical phenomena are the fundamentals of human learning also. They are, of course, much complicated in the more advanced states of human learning, such as the acquisition of skill with the violin, or of knowledge of the calculus, or of inventiveness in engineering. But it is impossible to understand the subtler and more planful learning of cultural men without clear ideas of the forces which make learning possible in its first form of directly connecting some gross bodily response with a situation immediately present to the senses.[3]

The mechanism for producing learning that Thorndike had hit upon in his "law of effect" was much more versatile than Pavlovian or classical conditioning. The latter is limited in that there has to be a preexisting reflex linking some natural stimulus and response; learning can only build upon this by "substituting" the conditioned stimulus for the natural one. But Thorndike's mechanism has no such limitation; any response to any stimulus can be conditioned simply by immediately reinforcing or rewarding it. The possibilities seem limitless, as was realized in the late 1930s by B. F. Skinner.

Case One

Before we look at Skinner's work, let us consider a simple case of one teacher giving some music instruction:

> Alright, play the chording as you practiced it.... You've lost the beat —keep time! OK, now try adding the right hand melody. Good. Wait, you've lost the timing again! Start again, please.

> OK... OK. Give it some feeling.... OK...OK...Good! Very good! Again....Now you've got it! One more time to be sure it's in your fingers now. Play, don't think about it. Play. Excellent! Next week we'll do the next movement. Try the same technique, first do the chording then add the melody.

Can you use Thorndike's theory to explain this bit of learning? Does Locke's or Plato's theory add anything? Go back to the case at the end of the last chapter in which Judith finally got the class to learn to be quiet after lunch. Does operant conditioning explain this case? Is it only the students who were learning—acquiring a new behavior? Is Judith, herself, like the cat learning to get out of the puzzle box? Are the students being reinforced? But if so, why aren't they being reinforced by the cookies?

B. F. Skinner

Working with laboratory rats and pigeons, Skinner made some remarkable advances. He discovered, again by serendipity, that an action or response does not have to be rewarded or reinforced every time it occurs; he found that his rats learned very effectively if they were rewarded fairly frequently but randomly, but even better, they then persisted with their newly learned behavior longer in the absence of reward. (An animal that has been rewarded for every action will soon stop performing if the rewards are cut off.) This is a common human experience; we all know that "a little praise goes a long way," and of course we find praise very rewarding. But we know also that we do not have to be rewarded for every right action we perform; an occasional pat on the back goes a long way. Teachers are well aware that praising a student for good work, or giving a gold star or some such reward, is very effective.

Skinner also found that he could "shape" the behavior of his laboratory animals in startling ways just by the judicious use of rewards. A famous demonstration, which he has done on film, is to teach pigeons to dance and to bowl a ball in a scaled-down ten-pin alley. Skinner takes a pigeon, and as soon as it makes the slightest move in the right direction it is rewarded. Then, when it eventually makes another move in the right direction, this is rewarded, and so it goes. In a matter of minutes Skinner can have the pigeon circling to the right or left. Quickly the dance is built up. He even has used this technique of "shaping behavior" via reinforcement to teach a pigeon to steer a guided missile!

He applied these principles to human learning, for example, by way of his work with the teaching machine and the programmed text. These devices work on the principle that humans find that getting the right answer

to a question is very rewarding. So in both the machine and the text, material to be learned is presented in small units, and the learner has to work through questions. If a correct answer is given, the learner receives immediate positive feedback and then moves on to the next item; if the answer is wrong, a "remedial" example or question is presented. As Skinner pointed out, the machine only presents material that the student is ready to learn (in the sense that the student got right all the steps that went before). The machine can give hints and prompts to help the student come up with the right answer, and "the machine, like the private tutor, reinforces the student for every correct response, using this immediate feedback not only to shape his behavior most efficiently but to maintain it in strength."[4]

Most people do not spend much time learning from teaching machines. But Skinner, no less than Thorndike, was confident that the mechanism of reinforcement of responses (operant behavior) was at work everywhere, in all types of learning. He wrote:

> While we are awake, we act upon the environment constantly, and many of the consequences of our actions are reinforcing. Through operant conditioning the environment builds the basic repertoire with which we keep our balance, walk, play games, handle instruments and tools, talk, write, sail a boat.[5]

To drive home the difference between Watson's behaviorism on one hand, and Thorndike's and Skinner's on the other, let us consider again the example of the child who is troublesome at bedtime. It will be recalled that Watson handled this case by locating some natural stimulus that produced the response (the unsettled behavior), and then he suggested that somehow the stimulus of turning off the light at bedtime had become substituted for this natural stimulus. Thorndike and Skinner have a different approach. They focus not on the stimulus that provokes the behavior, but instead they examine what happens to the behavior (or operant) after it has occurred. Thus, what happens to the child if he or she indulges in tantrums and so on at bedtime? Why, the parent comes in, turns on the light, and tries to comfort the child! In effect this is very rewarding—the child's response (the troublesomeness) becomes conditioned. No wonder it is repeated whenever the stimulus is the same, that is, at each bedtime. To change the behavior the parent has to extinguish this undesirable response, by ceasing to reward it; gradually, as was mentioned before, the now unrewarded behavior will die out.

Strengths and Weaknesses

Undoubtedly behaviorism has many positive features. For one thing, it is appealingly simple. It postulates that a single mechanism, conditioning, is responsible for producing learning, and furthermore, this mechanism operates throughout the whole animal kingdom. It is a mechanism that is easy for educators to master and put to good use; rewarding desirable behavior, and extinguishing (or even punishing) poor behavior, are techniques that all teachers can master (especially if they are rewarded for doing them). And there can be little doubt that these techniques are very effective—the animal experiments illustrate this, as do the variety of treatment programs that have been developed for human disorders such as autism, shyness, and antisocial behavior. The behaviorists, too, seem to bypass Plato's problem. Learning occurs not because there is something already present in the mind, but because, like other animals, we are "wired" so that any of our behavior that is reinforced is more likely to recur. Learning, then, can be described simply as the acquisition of new behavior without reference to mental events.

Skinner's writings not only contain entertaining expositions of all these things; he was also an able philosopher, and gave a powerful defense of the underlying assumptions. He believed that it is quite unscientific for psychologists to make use of the notions of "mind" and "consciousness" in their theories and explanations. Psychology, as a science, must only make use of data that are objective and publicly accessible. Operant behavior, and the consequences that this behavior produces, are things that meet this criterion, but what is happening in the mind of the learner does not meet it. Skinner did not deny that something might be happening in the "inner experience" of the learner, but unless this somehow connects with the observable realm it cannot be dealt with by science. Thus, he wrote:

> A purely private event would have no place in a study of behavior, or perhaps in any science; but events which are, for the moment at least, accessible only to the individual himself often occur as links in chains of otherwise public events and they must then be considered.[6]

All of this should force the educator to think seriously: Is learning best conceived as a change in behavior, or is something left out by this account? Are the events taking place in the mind of the learner of no relevance to the psychologist, and perhaps even more importantly, are they of no relevance to the work of the teacher? Just because operant conditioning undoubtedly has many useful applications in the classroom, does it mean that all cases of learning can be accounted for in these terms?

Has Plato's problem indeed been bypassed? Consider once more the learning of Einstein's theory. Maybe we can reinforce a learner after he or she correctly responds to a question about the theory, but it would be a tremendous fluke for the pupil to get it right in the first place unless the theory was understood, and to understand the theory wouldn't some

prior things have to be known? The ancient Greek is not so easy to dispose of! Talking about behavior certainly hides the problem, but it does not dispose of it.

Three more of the difficulties that have arisen for behaviorism are worth mentioning here. First, as mentioned earlier, Skinner's view of the nature of science was unnecessarily narrow. It is not true that scientists do not postulate unobservable entities or processes; in particle physics, for example, quarks with various quaint properties are supposed to exist, yet their connection to the actual data that physicists collect is extremely remote. Gravity is "invisible," but its effects are not. Second, Noam Chomsky has shown that a model such as Skinner's is not able, in principle, to account for the sort of phenomena that are met with in the field of linguistics. All of us, for example, can recognize the meaning of sentences and verbal constructions that we have never come across before. This is a particularly striking phenomenon in young children who are just learning their native language. So how could we have been reinforced to react to such sentences in some way, if we have never come across anything like them before? This cannot be how we learn our language. Finally, some experiments—of the sort that Skinner would admire—seem to indicate that something important is going on in the heads of learners, even laboratory rats. For instance, a rat might be taught to run a complex maze, by being positively reinforced by food when it reaches the target; then, if some of the passageways of the maze are closed off, or if the maze is rotated through a large angle, the rat can still get to the target using some new passage—showing, apparently, that it had formed some sort of "mental map" of the maze by which it could guide itself!

The suggestion in this last example is that learners may form some sort of mental structure of their knowledge, and it is this that allows them to perform or to behave correctly. There has been a great deal of research done on this in recent years, which we will discuss in chapters 5, 7, and 9. First, however, we must look at another piece of the "jigsaw puzzle" about learning, a contribution that was made by Gestalt researchers who, throughout much of the twentieth century, were rivals and critics of the behaviorists.

Case Two

To give you a feeling for the Gestalt theory of learning and to set you thinking about whether the operant conditioning theory works as easily as Skinner claims in the interpretation of all kinds of learning, consider the following case:

Joan had never before been interested in science. In fact, even though she was a bright student, she dreaded having to meet the middle school science requirements. But this term was different—really different. Ms. Cliff was an unusual teacher. Labs weren't done out of lab manuals. Instead, students arrived at the labs to find an array of items at each station and a "puzzle-problem" to deal with. This week the lab topic was electricity, which Ms. Cliff had talked about in class. On the lab table was a 6-volt battery that had a wire from each of its two terminals leading to two nails spaced about five inches apart and nailed into a board. One wire, on its way to its nail, was connected to a small flashlight bulb. The bulb was not lit. The problem-message for the lab was, "Try to make the bulb light using various materials available on the lab table to span the gap between the nails. Afterwards, write down what you think you learned from your experience." On the table was a pencil, a pen, a spoon, a piece of paper, a beaker of distilled water, a salt shaker, a piece of curly telephone cord, a piece of string, and a 50-cent piece. There was a little note tucked under the half dollar saying it had better be there when the lab was over!

Joan thought she might as well try each object in the order it was found on the table. She placed the pencil so it touched both nail heads. Nothing happened. The same nonresult occurred with the pen, but the spoon made the bulb glow as soon as it touched both nail heads. The paper didn't work and the beaker was a bit of a problem. Holding it so the glass bottom touched the nails produced nothing, but how, Joan thought, could she test the water in it? She thought and thought and then in a flash she realized she could pick up the board, turn it over, and submerge the nail heads in the water. But nothing happened ... maybe just a hint of a glow in the bulb, but not much. The salt shaker wouldn't reach the five inches, but Joan thought it might be interesting to shake some salt in the water to see if that made a difference. It did, the bulb glowed. She then poured a line of salt on the board to reach from one unit to the other. Nothing happened.

The piece of telephone cord was interesting. If she touched the outer plastic parts of the cord to the nails nothing happened, but if she touched the inner wire parts to the nails the bulb lit. The piece of string didn't work and Joan was left with the puzzle of how to test the half dollar. She thought and thought. There seemed to be no way to stretch it the five inches between the two nails and she knew it wouldn't dissolve in the water like the salt did. She just might not be able to test the half dollar. Already she suspected that metals were all good conductors of electricity and that if she could test the half dollar,

it would work. But how to do it? Then she "saw" the way! She quickly unwound the wire from the nails and placed the tip of each wire on opposite edges of the half dollar. It worked.

Now she had to think about and write down what she had learned.

Imagine that you are Joan and have just had this experience. What did you learn? Which of the theories dealt with thus far help explain any of the learnings?

Before going on to the next chapter, you may want to consider the argument in "Learning and Behavior Change," Chapter 10.